

UNITED STATES PATENT APPLICATION

FOR

METHOD AND SYSTEM FOR OPTIMIZING PRIVATE NETWORK FILE
TRANSFERS IN A PUBLIC PEER-TO-PEER NETWORK

Inventor(s):

Timothy S. DeBruine
Nicholas J. Hengeveld IV

Sawyer Law Group LLP
2465 E. Bayshore Road, Suite 406
Palo Alto, California 94303

METHOD AND SYSTEM FOR OPTIMIZING PRIVATE NETWORK FILE TRANSFERS IN A PUBLIC PEER-TO-PEER NETWORK

FIELD OF THE INVENTION

5 The present invention relates to peer-to-peer networks, and more particularly to a method and system for optimizing private network file transfers in a public peer-to-peer network.

BACKGROUND OF THE INVENTION

The Internet may be viewed as containing distributed information and centralized information. The distributed information is located throughout the Internet and typically takes the form of domain name servers and IP addresses, for instance. The centralized information is content, such as web pages and files, which is stored on and served by central servers.

15 Gaining access to such centralized content, however, is becoming increasingly difficult due to growing Internet congestion, limited bandwidth, and increasing file sizes (especially for media rich content). Traditional Internet technologies for distributing content, such as e-mail, streaming media, and FTP, have proven inadequate. E-mail is inadequate because due to the number of email messages and attachments passing through email servers, restrictions are placed
20 on the sizes of emails that restricts what can be sent as attachments. E-mail also

has security issues. PGP encryption is available for securing e-mails, but is not widely adopted.

5 Streaming media has the disadvantages of not working with all file types and is expensive because providers must purchase different software for the various streaming media standards. Streaming media also has not proven to be a reliable transfer method. And FTP file transfers also has disadvantages, which include being technically challenging to most users, causing configuration problems with firewalls, and suffering from inefficient file transfers. There are other solutions for distributing content, but they are usually proprietary and do not scale well.

15 Another problem with distributing centralized content is cost. As file sizes increase, the distribution of content is becoming increasingly expensive for content providers due to metered pricing of used bandwidth. In metered pricing, a content provider's Internet-Service-Provider (ISP) monitors the output of the servers used to provide the content, and charges the content provider 95% of the peak usage even though the average output is much lower. Thus, the cost of distributing content from central servers is one reason why attempts have been made to decentralize content.

20 One way to decentralize content is through peer-to-peer networks. Peer-to-peer network computing is a more efficient means for distributing resources and content over the Internet. In a peer-to-peer network, all workstations and computers

in the network may act as servers to all other users on the network. Some peer applications gain efficiencies by aggregating the distributed storage capacity of the computers across the network, such as Napster™ and Gnutella™, or aggregating the idle computing cycles of the computers, such as SETI@home™. Still others, such as instant messaging, take advantage of the direct network connections that peer devices can make to enhance communications.

Although peer networks are effective, current peer networks have inefficiencies. For example, it is not uncommon for a peer-to-peer network to have peers that are part of a private network, such as a local area network (LAN), for instance. When a peer requests a file from another peer, the file transfer typically occurs over the Internet, even when the two peers are within the same private network. In a peer-to-peer network that includes many private networks and many file transfers occurring within the same private network, transferring the file over the Internet is costly and wastes limited bandwidth.

Accordingly, what is needed is a an improved method and system for transferring files in a public peer-to-peer network when file transfers occur between computers belonging to a common private network. The present invention addresses such a need.

SUMMARY OF THE INVENTION

The present invention provides a method and system for optimizing private network file transfers in a public peer-to-peer network. The network includes a plurality of nodes wherein at least two of the nodes are part of the private network.

5 The method and system include receiving a search request from a first node for a file, and in response, determining that the file is stored on a second node. It is then determined whether the first and second nodes are part of the same private network, and if so, the second node is used to transfer the file to the first node over the private network, instead of the public network.

Accordingly, the present invention prevention spares network bandwidth minimizing the need to transfer the file over the Internet when a node is available to transfer the file that belongs to the same private network as the requesting node.

BRIEF DESCRIPTION OF THE DRAWINGS

Figures 1A and 1B are block diagrams illustrating a peer-to-peer (P2P) network architecture.

Figure 2 is a flow chart illustrating the process for registering a client node with the server node.

20 Figure 3 is a block diagram illustrating a preferred embodiment of the client application desktop window.

Figures 4A and 4B are flow charts illustrating the process of deciding whether a client node is locally reachable from the same private network as a requesting client node.

5 DETAILED DESCRIPTION

The present invention relates to facilitating file access on peer-to-peer networks. The following description is presented to enable one of ordinary skill in the art to make and use the invention and is provided in the context of a patent application and its requirements. Various modifications to the preferred embodiments will be readily apparent to those skilled in the art and the generic principles herein may be applied to other embodiments. Thus, the present invention is not intended to be limited to the embodiments shown but is to be accorded the widest scope consistent with the principles and features described herein.

Figures 1A and 1B are block diagrams illustrating a peer-to-peer (P2P) network architecture for use in accordance with one preferred embodiment of the present invention. The peer-to-peer network 10 includes a plurality of computers 18 interconnected over a network, such as Internet, where some of the computers 18 are configured as server nodes 12, and other computers 18 are configured as client nodes 14. A client node 14 may represent a single computer or a proprietary network, such as AOL, or a cable network, for example, and in a preferred embodiment, the server nodes 14 are located worldwide.

Any combination of server nodes 12 and client nodes 14 may form a private network 16, such as a local area network (LAN) or an extranet, which is a private network that uses the public Internet as its transmission system, but requires passwords to gain entrance. Some of the private networks 16 may be protected by a firewall 17. Firewalls 17 are widely used to give users of a private network 16 access to the Internet in a secure fashion as well as to separate a company's public web server from its internal network.

Figure 1B is a diagram illustrating contents of the server nodes 12 in a preferred embodiment of the present invention. Each server node 12 includes several databases for implementing the functions described above. The server node 12 includes a query database 24, a location database 26, a user database 28. The query and a location databases 24 and 26 store the names and locations of the files shared on the network, respectively. And the user database 32 includes account information for the users of the client nodes 14. In accordance with the present invention, the server node 12 also includes a node registry 30 for registering each node that is logged into the network 10.

The primary purpose of the peer-to-peer network 10 is the propagation of content files over the network 10. In a preferred embodiment, each server node 12 stores content 20 that comprises both commercial files 20a and noncommercial files 20b. Example type of content files may include audio MP3 files, video files, news articles and online magazines, image files, and confidential documents, for

instance. Once the content files have been downloaded from the server 12 to client nodes 14, the client nodes 14 serve the files directly to other client nodes 14. Thus, a need exist in the peer-to-peer network to allow each node 14 in the network 10 to share files with other nodes 14 in the network, regardless of whether two nodes 14 are separated by a firewall 17.

The present invention provides a method and system for optimizing private network file transfers in a public peer-to-peer network, such that nodes that are part of the same private network 16 share files by transferring the files within the private network 16, rather than transferring the files over the public network 10 (via the Internet). This is accomplished by recognizing when two nodes that need to transfer a file belong to the same private network, and causing the two nodes to send their request/responses to each other through their private network, rather than going through the Internet.

Figures 2-4 illustrate the process of optimizing private network file transfers in a public peer-to-peer network in accordance with a preferred embodiment of the present invention.

The optimization process begins with the registration of client nodes 14 with the server node 12, as illustrated in the flow chart of Figure 2. Once a user invokes the P2P client application 22 on their computer 18 in step 100, a TCP/IP connection is established with the server 12, and the client node determines its client IP

address in step 102. As is well known in the art, an IP (Internet Protocol) address is the address of a computer attached to a TCP/IP network. Every computer and server is assigned a unique IP address. Computers 18 have either a permanent address or one that is dynamically assigned to them each dial-up session.

5

During the TCP/IP connection, two interconnected nodes exchange requests/responses in the form of TCP/IP packets. Each TCP/IP packet sent out from the client node identifies the client node 14 and contains application data, the client IP address of the client node 14, and a destination IP address of the recipient. IP addresses are written as four sets of numbers separated by periods; for example, 204.171.64.2, that includes a network address (netid) that identifies the private network or subnet, and a host address (hostid) that identifies the computer within the private network. The subnet mask is a method used for splitting IP networks into a series of subgroups, or subnets. The mask is a binary pattern that is matched up with the IP address to turn part of the host ID address field into a field for subnets. Computers 18 can determine their subnet mask by making an operating system call.

When a client node 14 is part of a private network, the private network may or may not be protected by a firewall, and the firewall may or may not perform network address translation (NAT). NAT is a process whereby a firewall translates a range of IP addresses to another as packets are routed between networks. NAT

also keeps individual addresses of the private network hidden from the outside world.

5 After the client node 14 establishes a TCP/IP connection, the server 12 obtains the client IP address and the subnet mask from the TCP/IP packet, and determines the peer IP address of the client node 14 by sending a probe message to the client node 14 and observing what IP address the client node 14 is connecting from in step 104. Thereafter, a registration process is initiated whereby the client node 14's peer IP address, client IP address, and subnet mask node are registered in the node registry 30 in step 106.

10 The node registry 30 then determines if the client node 14 is protected by a firewall by comparing the observed peer IP address with the reported client IP address in step 108. If the peer IP address does match the client IP address, then the entry for the client node 14 in the node registry 30 is set to indicate that both NAT has been performed on the client node 14, and that the client node 14 is unreachable from the network 10 in step 110.

15 If the peer IP address matches the client IP address, then the entry for the client node 14 is set to indicate that NAT has not been performed on the client node 14 in step 112. If the server can connect to the client node 14 through the client IP address in step 114, then the entry for the client node 14 in the node registry 30 is set to indicate that the client node 14 is directly reachable from the network 10 in

step 116. If the server node 12 cannot connect to the client node 14, then the node registry is set to indicate that the node is unreachable from the network 10 in step 118.

5 Besides registering with the node registry 30, the client application 22 also displays a client application desktop window on the computer to allow the user to share files on the network 10. Referring now to Figure 3, a block diagram is shown illustrating a preferred embodiment of the client application desktop window. The client application 22 allows the user to perform three primary functions: publish over the network, receive files over the network, and search for files to download. The client application desktop window 50 may include a row of command buttons 52, and an area 54 for displaying folders and icons. The user logs in and out of the network 10 via command buttons 52a and 52b, and may search for files on the network via the search button 52c. An inbox folder 56 contains files that are received over the network 10, and a shared folder 58 contains files that the user wishes to publish over the network 10 for access by other client nodes 14. User icons 60 represent individual users and groups of users to which the user wants to exchange files with on a peer-to-peer basis.

20 Figures 4A and 4B are flow charts illustrating the process of deciding whether a client node is locally reachable from the same private network as a requesting client node in response to the user clicking on the search button 52c and entering search terms. Once the server node 12 receives the search terms, it is

5 determined if there are any nodes in the network that contain a file matching the search terms in step 150. This is accomplished by searching the query database 24 for file names that match the search terms and by then finding the nodes containing that file by querying the location database 26. The client nodes 14 containing the file will hereinafter be referred to as target nodes.

10 According to the present invention, instead of just displaying a list of matching file names on all nodes throughout the network 10, the server node 12 further determines whether there are any target nodes within the same private network as the requesting node by looking up the requesting node and the target node in the node registry 30 in step 152. If the node registry 30 indicates that NAT has been performed on both nodes and at the peer IDs of both nodes match in step 154, then the target node is considered to be within the same network as the requesting client node and therefore "locally reachable" through its local client IP address from the requesting client node in step 158.

15 If the condition of step 152 fails, but the node registry 30 indicates that NAT has not been performed on either of the client nodes 14 and that the subnet IDs of each of the client nodes 14 match in step 156, then the target node is still
20 considered to be within the same network as the requesting client node and locally reachable from the requesting client node in step 158. If the target node is found to be locally reachable, then it is added to the search results list that will be returned to the requesting client node in step 160.

15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
100
105
110
115
120
125
130
135
140
145
150
155
160
165
170
175
180
185
190
195
200
205
210
215
220
225
230
235
240
245
250
255
260
265
270
275
280
285
290
295
300
305
310
315
320
325
330
335
340
345
350
355
360
365
370
375
380
385
390
395
400
405
410
415
420
425
430
435
440
445
450
455
460
465
470
475
480
485
490
495
500
505
510
515
520
525
530
535
540
545
550
555
560
565
570
575
580
585
590
595
600
605
610
615
620
625
630
635
640
645
650
655
660
665
670
675
680
685
690
695
700
705
710
715
720
725
730
735
740
745
750
755
760
765
770
775
780
785
790
795
800
805
810
815
820
825
830
835
840
845
850
855
860
865
870
875
880
885
890
895
900
905
910
915
920
925
930
935
940
945
950
955
960
965
970
975
980
985
990
995

If the target node is not locally reachable (not within the same private network as the requesting node), then it is determined whether the requesting node can communicate directly with the target node by looking up the requesting and target nodes in the node registry 30 to determine if it is registered as being directly reachable through its peer ID in step 162. If the target node is directly reachable, the target node is added to the list of search results in step 164. If none of the conditions above are not met, then it is determined that the client node 14 containing file is not reachable from the requesting node.

Referring now to Figure 4B, after all the target nodes are examined, the search results are sorted first by locally reachable client nodes 14 followed by the directly reachable client nodes 14 and then returned to the requesting node in step 164. In a preferred embodiment, the name of the file, and the identities and addresses of the client node(s) 14 are included in the search result sent to the requesting node.

After the search results are received on the requesting node, the user may click on the file name displayed next to an identified node to obtain the file in step 166. If the target node is locally reachable, then the client application 22 sends the request for the file to the target node using the client IP address of the target node in step 168. If the target node is directly reachable, then the client application sends the request for the file to the target node using the peer IP address in step 170. The

target node 14 then responds by sending the file to the requesting node using the requesting node's client IP address in step 172.

Thus, when the target node is within the same private network as the requesting node (locally reachable), the present invention enables the two nodes to communicate such that the file is transferred over the private network.

In an alternative embodiment, the server node 12 may return a list of search results to the requesting node, where the list includes the identities and addresses of the matching nodes, their IP addresses and subnet masks, port, reachability status and so on, to the requesting node. The requesting node may then perform the comparisons described above to determine if any of the nodes are on the same private network.

According to the present invention, file transfers on the P2P network 10 are optimized by identifying when a node requesting a file is on the same private network as a second node containing the file, and using the second node to transfer the file using over the private network. By not sending the file from another node that is not part of the private network, an Internet file transfer does not take place, sparing network bandwidth.

A method and system for optimizing private network file transfers in a public peer-to-peer network has been disclosed. Although the present invention has been

